



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[session* and (time* or length* or duration*)<AND>((sessions<AND>((packet and relay <AND>((load and balance* and server and service))))))]**

Found **43** of **127,944** searched.

Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 20 of 43 **short listing**



1

2

3



- 1 Fast detection of communication patterns in distributed executions 87%



Thomas Kunz , Michiel F. H. Seuren

Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research November 1997

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

- 2 Network support for mobile multimedia using a self-adaptive distributed proxy 83%



Zhuoqing Morley Mao , Hoi-sheung Wilson So , Byunghoon Kang

Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video January 2001

Recent advancements in video and audio codec technologies~(e.g., RealV ideo [18] make multimedia streaming possible across a wide range of network conditions. With an increasing trend of ubiquitous connectivity, more and more areas have overlapping coverage of multiple wired and wireless networks. Because the best network service changes as the user moves, to provide good multimedia application performance, the service needs to adapt to user movement as well as network and computational res ...

- 3 Special feature: Report on a working session on security in wireless ad hoc networks 82%



Levente Buttyán , Jean-Pierre Hubaux

ACM SIGMOBILE Mobile Computing and Communications Review January 2003
Volume 7 Issue 1

- 4** Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer 82%



Michael J. Freedman , Robert Morris

Proceedings of the 9th ACM conference on Computer and communications security November 2002

Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a restricted topology in a way that adversaries cannot easily influence ...

- 5** Using name-based mappings to increase hit rates 82%



David G. Thaler , Chinya V. Ravishankar

IEEE/ACM Transactions on Networking (TON) February 1998
Volume 6 Issue 1

- 6** Design and modelling of internode: a mobile provider provisioned VPN 82%



Francisco Barceló , Josep Paradells , Fofy Setaki , Monique Gibeaux

Mobile Networks and Applications February 2003

Volume 8 Issue 1

This paper presents the design and architecture of a mobile Provider Provisioned VPN (PPVPN) together with a performance evaluation oriented model that allows first estimates of the VPN set-up delay to be computed. At the same time, some consequences of the discussion can be applied to the design of the VPN configuration parameters. Many different technologies and protocols are used: access is supplied through GPRS or WaveLANs, IP mobility is supported by Mobile IP, and the VPN is based on the I ...

- 7** A peering architecture for ubiquitous IP multicast streaming 80%



C. K. Yeo , B. S. Lee , M. H. Er

ACM SIGOPS Operating Systems Review July 2002

Volume 36 Issue 3

IP Multicast has been widely used for the distribution of real-time video. However, a large portion of the Internet is not multicast-enabled. Such networks rely on static relays and static tunnels to support video streaming. Being static, the streaming cannot readily adjust to the changing network conditions. This paper proposes an application level overlay using peering technology for seamless ubiquitous multimedia streaming across both multicast and unicast networks. The overlay comprises an a ...

- 8** Session A: Routing: On the impact of alternate path routing for load balancing in mobile ad hoc networks 80%



Marc R. Pearlman , Zygmunt J. Haas , Peter Sholander , Siamak S. Tabrizi

Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing November 2000

Alternate path routing (APR) can provide load balancing and route failure protection by distributing traffic among a set of diverse paths. These benefits make APR appear to be an ideal candidate for the bandwidth limited and mobile ad-hoc networks. However,

we find that APR's potential is not fully realized in ad-hoc networks because of route coupling resulting from the geographic proximity of candidate paths between common endpoints. In multiple channel networks, coupling occurs when paths shar ...

9 Session summaries from the 17th symposium on operating systems 80%

 principle (SOSP'99)

Jay Lepreau , Eric Eide

ACM SIGOPS Operating Systems Review April 2000

Volume 34 Issue 2

10 Increasing the realibility of email services 80%

 Joe Armstrong

Proceedings of the 2000 ACM symposium on Applied computing March 2000

11 A mechanism for supporting client migration in a shared window system 80%

 Goopeel Chung , Prasun Dewan

Proceedings of the 9th annual ACM symposium on User interface software and technology November 1996

12 Stateful distributed interposition 80%

 John Reumann , Kang G. Shin

ACM Transactions on Computer Systems (TOCS) February 2004

Volume 22 Issue 1

Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex servic ...

13 Application level performance: DNS performance and the effectiveness 80%

 of caching

Jaeyeon Jung , Emil Sit , Hari Balakrishnan , Robert Morris

Proceedings of the First ACM SIGCOMM Workshop on Internet Measurement Workshop November 2001

This paper presents a detailed analysis of traces of DNS and associated TCP traffic collected on the Internet links of the MIT Laboratory for Computer Science and the Korea Advanced Institute of Science and Technology (KAIST). The first part of the analysis details how clients at these institutions interact with the wide-area DNS system, focusing on performance and prevalence of failures. The second part evaluates the effectiveness of DNS caching. In the most recent MIT trace, 23% of lookups rece ...

14 A measurement analysis of Internet traffic over frame relay 80%

 Judith L. Jerkins , John Monroe , Jonathan L. Wang

ACM SIGMETRICS Performance Evaluation Review September 1999

Volume 27 Issue 2

Various approaches have been proposed and implemented to relieve the congestion in the Public Switched Telephone Networks (PSTNs) induced by recent meteoric growth of Internet services. The Internet/Intranet Transport Service (IITS) offered by the

Southwestern Bell Telephone (SWBT) Company provides an example of one such implementation which off-loads the long-holding time data traffic from PSTNs to a packet technology (in this case Frame Relay). This paper describes analysis of 1997 IITS traffic ...

15 Manufacturing resource planning on a PC local area network 80%


 H. Clark Kee , Roy L. Post
ACM SIGAPL APL Quote Quad , Proceedings of the international conference on APL May 1986
 Volume 16 Issue 4

This paper details a large APL programming project of 12 man years. An integrated software system structured on the principles of MRP (manufacturing resource planning) was implemented by a Bristol-Myers in house team for use in a new manufacturing facility. The system applies off-the-shelf technology in innovative ways, using STSC APL*PLUS/PC as the only programming language, to build a very sophisticated application on IBM/PCs fully sharing data in a secure environment via the N ...

16 Protecting web servers from distributed denial of service attacks 77%

 Frank Kargl , Joern Maier , Michael Weber
Proceedings of the tenth international conference on World Wide Web April 2001

17 A multi-user framework supporting video-based avatars 77%

 Peter Quax , Tom Jehaes , Pieter Jorissen , Wim Lamotte
Proceedings of the 2nd workshop on Network and system support for games May 2003


In this paper we present our ongoing work in setting up a multi-user framework that supports video texturing on avatars, creating added value for both gaming and collaborative work applications. Limiting the data propagation and bandwidth usage is a key goal when targeting a scalable application that will be deployed on a general-purpose network such as the Internet. We therefore present a number of techniques that can be used in these circumstances, including increased client responsibilities a ...

18 Session 1: QoS's downfall: at the bottom, or not at all! 77%

 Jon Crowcroft , Steven Hand , Richard Mortier , Timothy Roscoe , Andrew Warfield
Proceedings of the ACM SIGCOMM workshop on Revisiting IP QoS: What have we learned, why do we care? August 2003

Quality of Service (QoS) has been touted as a technological requirement for many different networks at many different times. However, very few (if any) schemes for providing it have ever been successful, despite a huge amount of research in the area of QoS provision. In this position paper we analyze some of the reasons why so many QoS mechanisms have failed to be widely deployed. We suggest two factors in this failure: the timeliness of QoS mechanisms (they rarely arrive when they are needed), ...

19 DOS protection: Using graphic turing tests to counter automated DDoS attacks against web servers 77%

 William G. Morein , Angelos Stavrou , Debra L. Cook , Angelos D. Keromytis , Vishal Misra , Dan Rubenstein
Proceedings of the 10th ACM conference on Computer and communication security October 2003

We present WebSOS, a novel overlay-based architecture that provides guaranteed

access to a web server that is targeted by a denial of service (DoS) attack. Our approach exploits two key characteristics of the web environment: its design around a human-centric interface, and the extensibility inherent in many browsers through downloadable "applets." We guarantee access to a web server for a large number of *previously unknown* users, without requiring pre-existing trust relationships between ...

20 Overlay & peer-to-peer networks: SplitStream: high-bandwidth

77%



multicast in cooperative environments

Miguel Castro , Peter Druschel , Anne-Marie Kermarrec , Animesh Nandi , Antony Rowstron , Atul Singh

Proceedings of the nineteenth ACM symposium on Operating systems principles
October 2003

In tree-based multicast systems, a relatively small number of interior nodes carry the load of forwarding multicast messages. This works well when the interior nodes are highly-available, dedicated infrastructure routers but it poses a problem for application-level multicast in peer-to-peer systems. SplitStream addresses this problem by striping the content across a forest of interior-node-disjoint multicast trees that distributes the forwarding load among all participating peers. For example, i ...

Results 1 - 20 of 43 short listing

 **Prev**
Page 1 2 3  **Next**
Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[(simulat* or emulat* or model*)<AND>(((mean or median or average)<AND>((session* and (time* or length* or duration*)<AND>((sessions<AND>((packet and relay <AND>((load and balance* and server and service)))))))))]]**

Found **43** of **127,944** searched.

Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: Title Publication Publication Date Score Binder

Results 21 - 40 of 43 short listing



1 2 3



21 Session II: Paging and performance in cellular networks: P-MIP: paging 77%



in mobile IP

Xiaowei Zhang , Javier Gomez Castellanos , Andrew T. Campbell

Proceedings of the 4th ACM international workshop on Wireless mobile multimedia July 2001

As the number of Mobile IP users grows, so will the signaling overhead associated with Internet mobility management in the core IP network. This presents a significant challenge to Mobile IP as the number of mobile devices scale-up. In cellular networks, registration and paging techniques are used to minimize the signaling overhead and optimize the mobility management performance. Currently, Mobile IP supports registration but not paging. In this paper, we argue that Mobile IP should be extended ...

22 Best poster papers from MobiHoc 2002: Virtual operator based AAA in 77%








wireless LAN hot spots with ad-hoc networking support

Junbiao Zhang , Jun Li , Stephen Weinstein , Nan Tu


ACM SIGMOBILE Mobile Computing and Communications Review June 2002
Volume 6 Issue 3

Sound and effective authentication, authorization and accounting (AAA) schemes for convenient and secure mobile wireless accesses are of great importance given the increased popularity and business opportunities in public wireless LAN hot spots. One possible scheme, which uses the mobile users' service providers as the single point of contact for all AAA transactions, is emerging as a very promising solution. We refer to such service providers as "virtual operators". In this paper, we discuss va ...

- 23** A framework for the transmission of streaming media to mobile devices 77%
 Kevin Curran , Gerard Parr
International Journal of Network Management January 2002
Volume 12 Issue 1
One interesting problem is the delay imposed upon mobile receivers when switching between wireless cells. We provide a solution to this in the form of an extension of Mobile IP's handoff algorithm. Our solution involves the exploitation of mobility prediction to predict a mobile terminal's future location based on its previous history (i.e. the last cell that it has been in) and for the media stream to be already present and cached by next cells base station ready for receiving by the mobile dev ...
- 24** Network infrastructure for massively distributed games 77%
 Daniel Bauer , Sean Rooney , Paolo Scotton
Proceedings of the 1st workshop on Network and system support for games April 2002
The popularity of hypertext documents led to the need for specific network infrastructure elements such as HTML caches, URL-based switches, web-server farms, and as a result created several new industries as companies rushed to fill that need. We contend that massive distributed games will have a similar impact on the Internet and will require similar dedicated support. This paper outlines some initial work on prototyping such support. Our approach is to combine highlevel game specific logic and ...
- 25** Poster Session 2: Power-aware source routing protocol for mobile ad hoc networks 77%
 Morteza Maleki , Karthik Dantu , Massoud Pedram
Proceedings of the 2002 international symposium on Low power electronics and design August 2002
Ad hoc wireless networks are power constrained since nodes operate with limited battery energy. To maximize the lifetime of these networks (defined by the condition that a fixed percentage of the nodes in the network "die out" due to lack of energy), network-related transactions through each mobile node must be controlled such that the power dissipation rates of all nodes are nearly the same. Assuming that all nodes start with a finite amount of battery capacity and that the energy dissipation p ...
- 26** Virtual environments and HLA: A scalable architecture for supporting interactive games on the internet 77%
 Wentong Cai , Percival Xavier , Stephen J. Turner , Bu-Sung Lee
Proceedings of the sixteenth workshop on Parallel and distributed simulation May 2002
This paper presents a scalable architecture for supporting large-scale interactive Internet games. In order to support a large number of participants and to divide the workload, the virtual world is divided into partitions. Each partition is then assigned to a server. A client (i.e., a player or a participant) will join a server according to the position of the avatar it controls. Compared to a centralized architecture, this distributed client-server architecture is more scalable. In addition, c ...
- 27** The state of the art in locally distributed Web-server systems 77%
 Valeria Cardellini , Emiliano Casalicchio , Michele Colajanni , Philip S. Yu
ACM Computing Surveys (CSUR) June 2002
Volume 34 Issue 2
The overall increase in traffic on the World Wide Web is augmenting user-perceived response times from popular Web sites, especially in conjunction with special events.


System platforms that do not replicate information content cannot provide the needed scalability to handle large traffic volumes and to match rapid and dramatic changes in the number of clients. The need to improve the performance of Web-based services has produced a variety of novel content delivery architectures. This article w ...

28 An architecture for secure wide-area service discovery 77%

 Todd D. Hodes , Steven E. Czerwinski , Ben Y. Zhao , Anthony D. Joseph , Randy H. Katz
Wireless Networks March 2002
 Volume 8 Issue 2/3

The widespread deployment of inexpensive communications technology, computational resources in the networking infrastructure, and network-enabled end devices poses an interesting problem for end users: how to locate a particular network service or device out of hundreds of thousands of accessible services and devices. This paper presents the architecture and implementation of a secure wide-area Service Discovery Service (SDS). Service providers use the SDS to advertise descriptions of available ...

29 P-MIP: paging extensions for mobile IP 77%

 Xiaowei Zhang , Javier Gomez Castellanos , Andrew T. Campbell
Mobile Networks and Applications April 2002
 Volume 7 Issue 2


As the number of Mobile IP users grows, so will the signalling overhead associated with Internet mobility management in the core IP network. This presents a significant challenge to Mobile IP as the number of mobile devices scale-up. In cellular networks, registration and paging techniques are used to minimize the signalling overhead and optimize the mobility management performance. Currently, Mobile IP supports registration but not paging. In this paper, we argue that Mobile IP should be extend ...

30 Applications: YouServ: a web-hosting and content sharing tool for the masses 77%

 Roberto J. Bayardo Jr. , Rakesh Agrawal , Daniel Gruhl , Amit Somani
Proceedings of the eleventh international conference on World Wide Web May 2002

YouServ is a system that allows its users to pool existing desktop computing resources for *high availability* web hosting and file sharing. By exploiting standard web and internet protocols (e.g. HTTP and DNS), YouServ does not require those who access YouServ-published content to install special purpose software. Because it requires minimal server-side resources and administration, YouServ can be provided at a very low cost. We describe the design, implementation, and a successful intrane ...

31 Limits of wide-area thin-client computing 77%

 Albert Lai , Jason Nieh
ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems June 2002
 Volume 30 Issue 1

While many application service providers have proposed using thin-client computing to deliver computational services over the Internet, little work has been done to evaluate the effectiveness of thin-client computing in a wide-area network. To assess the potential of thin-client computing in the context of future commodity high-bandwidth Internet access, we have used a novel, non-invasive slow-motion benchmarking technique to evaluate the performance of several popular thin-client computing plat ...

77%

32 Session 1: Applications: Convenient abstractions in stormcast applications

Dag Johansen , Gunnar Hartvigsen

Proceedings of the 6th workshop on ACM SIGOPS European workshop: Matching operating systems to application needs September 1994

In this paper we present experience with meteorology applications and appropriate distributed computing abstractions. We focus on the need for co-existence and integration of multiple paradigms in large scale distributed applications, rather than enforcing a favourite paradigm whenever possible.

33 Network support for mobile multimedia using a self-adaptive distributed 77% proxy

Zhuoqing Morley Mao , Hoi-sheung Wilson So , Byunghoon Kang

Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video January 2001

Recent advancements in video and audio codec technologies~(e.g., RealV ideo [18] make multimedia streaming possible across a wide range of network conditions. With an increasing trend of ubiquitous connectivity, more and more areas have overlapping coverage of multiple wired and wireless networks. Because the best network service changes as the user moves, to provide good multimedia application performance, the service needs to adapt to user movement as well as network and computational res ...

34 Dimensioning server access bandwidth and multicast routing in overlay 77% networks

Sherlia Y. Shi , Jonathan S. Turner , Marcel Waldvogel

Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video January 2001

Application-level multicast is a new mechanism for enabling multicast in the Internet. Driven by the fast growth of network audio/video streams, application-level multicast has become increasingly important for its efficiency of data delivery and its ability of providing value-added services to satisfy application specific requirements. From a network design perspective, application-level multicast differs drastically from traditional IP multicast in its network cost model and routing stra ...

35 FIRE: flexible Intra-AS routing environment

77%

 Craig Partridge , Alex C. Snoeren , W. Timothy Strayer , Beverly Schwartz , Matthew Condell , Isidro Castiñeyra**ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication** August 2000

Volume 30 Issue 4

Current routing protocols are monolithic, specifying the algorithm used to construct forwarding tables, the metric used by the algorithm (generally some form of hop-count), and the protocol used to distribute these metrics as an integrated package. The Flexible Intra-AS Routing Environment (FIRE) is a link-state, intra-domain routing protocol that decouples these components. FIRE supports run-time-programmable algorithms and metrics over a secure link-state distribution protocol. By allow ...


36 Increasing the realibility of email services

77%

 Joe Armstrong**Proceedings of the 2000 ACM symposium on Applied computing** March 2000

37 Bimodal multicast

77%


 Kenneth P. Birman , Mark Hayden , Oznur Ozkasap , Zhen Xiao , Mihai Budiu , Yaron Minsky**ACM Transactions on Computer Systems (TOCS)** May 1999

Volume 17 Issue 2


There are many methods for making a multicast protocol "reliable." At one end of the spectrum, a reliable multicast protocol might offer toticity guarantees, such as all-or-nothing delivery, delivery ordering, and perhaps additional properties such as virtually synchronous addressing. At the other are protocols that use local repair to overcome transient packet loss in the network, offering "best effort" reliability. Yet none of this prior work has treated stability ...

38 How can routers help Internet economics?

77%

 John M. Schnizlein**Proceedings of the first international conference on Information and computation economies** October 1998**39 Mobile networking in the Internet**

77%


 Charles E. Perkins**Mobile Networks and Applications** December 1998

Volume 3 Issue 4

Computers capable of attaching to the Internet from many places are likely to grow in popularity until they dominate the population of the Internet. Consequently, protocol research has shifted into high gear to develop appropriate network protocols for supporting mobility. This introductory article attempts to outline some of the many promising and interesting research directions. The papers in this special issue indicate the diversity of viewpoints within the research community, and it is ...

40 Adapting to network and client variability via on-demand dynamic

77%

 distillation

Armando Fox , Steven D. Gribble , Eric A. Brewer , Elan Amir

Proceedings of the seventh international conference on Architectural support for programming languages and operating systems October 1996

Volume 30 , 31 Issue 5 , 9

The explosive growth of the Internet and the proliferation of smart cellular phones and handheld wireless devices is widening an already large gap between Internet clients. Clients vary in their hardware resources, software sophistication, and quality of connectivity, yet server support for client variation ranges from relatively poor to none at all. In this paper we introduce some design principles that we believe are fundamental to providing "meaningful" Internet access for the entire range of ...

Results 21 - 40 of 43 **short listing**
Prev
Page**1****2****3**
Next
Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[(simulat* or emulat* or model*)<AND>(((mean or median or average)<AND>((session* and (time* or length* or duration*)<AND>((sessions<AND>((packet and relay <AND>((load and balance* and server and service)))))))))]]**

Found **43** of **127,944** searched.

Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results **41 - 43** of **43** short listing



1

2

3



41 A mechanism for supporting client migration in a shared window system 77%



Goopeel Chung , Prasun Dewan

Proceedings of the 9th annual ACM symposium on User interface software and technology November 1996

42 Towards an active network architecture 77%



David L. Tennenhouse , David J. Wetherall

ACM SIGCOMM Computer Communication Review April 1996

Volume 26 Issue 2

Active networks allow their users to inject customized programs into the nodes of the network. An extreme case, in which we are most interested, replaces packets with "capsules" - program fragments that are executed at each network router/switch they traverse. Active architectures permit a massive increase in the sophistication of the computation that is performed within the network. They will enable new applications, especially those based on application-specific multicast, information fusion, a ...

43 Credit-based flow control for ATM networks: credit update protocol, 77%



adaptive credit allocation and statistical multiplexing

H. T. Kung , Trevor Blackwell , Alan Chapman

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Communications architectures, protocols and applications October 1994

Volume 24 Issue 4

This paper presents three new results concerning credit-based flow control for ATM networks: (1) a simple and robust credit update protocol (CUP) suited for relatively inexpensive hardware/software implementation; (2) automatic adaptation of credit

buffer allocation for virtual circuits (VCs) sharing the same buffer pool; (3) use of credit-based flow control to improve the effectiveness of statistical multiplexing in minimizing switch memory. These results have been substantiated by analysi ...

Results 41 - 43 of 43 **short listing**


Prev
Page

1

2

3


Next
Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[load and balance* and server and service]**

Found **2,059** of **127,944** searched.

Warning: Maximum result set of 200 exceeded. Consider refining.

Search within Results

GO

[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: Title Publication Publication Date Score

Results 1 - 20 of 200

short listing

Prev
Page

1 2 3 4 5 6 7 8 9 10

Next
Page

- 1** Optimizing data aggregation for cluster-based internet services 100%

Lingkun Chu , Hong Tang , Tao Yang , Kai Shen
ACM SIGPLAN Notices , Proceedings of the ninth ACM SIGPLAN symposium on Principles and practice of parallel programming June 2003
 Volume 38 Issue 10
 Large-scale cluster-based Internet services often host partitioned datasets to provide incremental scalability. The aggregation of results produced from multiple partitions is a fundamental building block for the delivery of these services. This paper presents the design and implementation of a programming primitive -- Data Aggregation Call (DAC) -- to exploit partition parallelism for cluster-based Internet services. A DAC request specifies a local processing operator and a global reduction ope ...
- 2** The state of the art in locally distributed Web-server systems 100%

Valeria Cardellini , Emiliano Casalicchio , Michele Colajanni , Philip S. Yu
ACM Computing Surveys (CSUR) June 2002
 Volume 34 Issue 2
 The overall increase in traffic on the World Wide Web is augmenting user-perceived response times from popular Web sites, especially in conjunction with special events. System platforms that do not replicate information content cannot provide the needed scalability to handle large traffic volumes and to match rapid and dramatic changes in the number of clients. The need to improve the performance of Web-based services has produced a variety of novel content delivery architectures. This article w ...
- 3** SEDA: an architecture for well-conditioned, scalable internet services 100%

Matt Welsh , David Culler , Eric Brewer
ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM

symposium on Operating systems principles October 2001

Volume 35 Issue 5

We propose a new design for highly concurrent Internet services, which we call the *staged event-driven architecture* (SEDA). SEDA is intended to support massive concurrency demands and simplify the construction of well-conditioned services. In SEDA, applications consist of a network of event-driven *stages* connected by explicit *queues*. This architecture allows services to be well-conditioned to load, preventing resources from being overcommitted when demand exceeds service cap ...

4 Managing energy and server resources in hosting centers 100%

Jeffrey S. Chase , Darrell C. Anderson , Prachi N. Thakar , Amin M. Vahdat , Ronald P. Doyle

ACM SIGOPS Operating Systems Review , Proceedings of the eighteenth ACM symposium on Operating systems principles October 2001

Volume 35 Issue 5

Internet hosting centers serve multiple service sites from a common hardware base. This paper presents the design and implementation of an architecture for resource management in a hosting center operating system, with an emphasis on *energy* as a driving resource management issue for large server clusters. The goals are to provision server resources for co-hosted services in a way that automatically adapts to offered load, improve the energy efficiency of server clusters by dynamically res ...

5 Parallel and distributed systems and networking: Load balancing for 100%

the management of service performance in open service markets: a customer-oriented approach

Dirk Thißen

Proceedings of the 2002 ACM symposium on Applied computing March 2002

In open service markets customers can choose between several providers offering similar services. To survive in the arising competition, service providers are compelled to satisfy their customers by not only offering services for a reasonable price but additionally deploy them in an efficient way regarding e.g. performance and availability. Because in a service provision process application, network, and system aspects are involved, new management concepts are needed. This paper discusses a mech ...

6 Architectural considerations for next generation file systems 100%

Prashant Shenoy , Pawan Goyal , Harrick M. Vin

Proceedings of the seventh ACM international conference on Multimedia (Part 1) October 1999

We evaluate two architectural alternatives—partitioned and integrated—for designing next generation file systems. Whereas a partitioned server employs a separate file system for each application class, an integrated file server multiplexes its resources among all application classes; we evaluate the performance of the two architectures with respect to sharing of disk bandwidth among the application classes. We show that although the problem of sharing disk bandwidth in integrate ...

7 A load cluster management system using SNMP and web 100%

Myung-Sup Kim , Mi-Joung Choi , James W. Hong

International Journal of Network Management November 2002

Volume 12 Issue 6

Clustered servers for Internet service is a popular solution to cope with the explosive increase in client requests. The high probability of service failure in cluster servers make the cluster management system necessary to provide high availability and

convenient administrator control. In this paper, we present the design and implementation of a load cluster management system (LCMS) based on SNMP and Web technology. Our LCMS implementation has been deployed on a commercial ultra-dense server.

- 8** A client-aware dispatching algorithm for web clusters providing multiple services 100%



Emiliano Casalicchio , Michele Colajanni

Proceedings of the tenth international conference on World Wide Web April 2001

- 9** Distributed operating systems 100%



Andrew S. Tanenbaum , Robbert Van Renesse

ACM Computing Surveys (CSUR) December 1985

Volume 17 Issue 4

Distributed operating systems have many aspects in common with centralized ones, but they also differ in certain ways. This paper is intended as an introduction to distributed operating systems, and especially to current university research about them. After a discussion of what constitutes a distributed operating system and how it is distinguished from a computer network, various key design issues are discussed. Then several examples of current research projects are examined in some detail ...

- 10** Protecting web servers from distributed denial of service attacks 100%



Frank Kargl , Joern Maier , Michael Weber

Proceedings of the tenth international conference on World Wide Web April 2001

- 11** Efficient support for interactive service in multi-resolution VOD systems 100%



Kelvin K. W. Law , John C. S. Lui , Leana Golubchik

The VLDB Journal — The International Journal on Very Large Data Bases October 1999

Volume 8 Issue 2

Advances in high-speed networks and multimedia technologies have made it feasible to provide video-on-demand (VOD) services to users. However, it is still a challenging task to design a cost-effective VOD system that can support a large number of clients (who may have different quality of service (QoS) requirements) and, at the same time, provide different types of VCR functionalities. Although it has been recognized that VCR operations are important functionalities in providing VOD service, tec ...

- 12** Issues in the Design of Adaptive Middleware Load Balancing 100%



Ossama Othman , Douglas C. Schmidt

ACM SIGPLAN Notices August 2001

Volume 36 Issue 8

Load balancing middleware is used extensively to improve scalability and overall system throughput in distributed systems. Many load balancing middleware services are simplistic, however, since they are geared only for specic use-cases and environments. These limitations make it hard to use the same load balancing service for anything other than the distributed application it was designed for originally. This lack of generality forces continuous re-development of application-specic load balancin ...

- 13** The aggregate server method for analyzing serialization delays in computer systems 100%



Subhash C. Agrawal , Jeffrey P. Buzen
ACM Transactions on Computer Systems (TOCS) May 1983
 Volume 1 Issue 2

14 Design and implementation of a portable and adaptable load balancing framework 100%



Erik Putrycz

**Proceedings of the 2003 conference of the Centre for Advanced Studies
 conference on Collaborative research** October 2003

Scaling applications to large networks and an increasing number of users has been since years a technical challenge. Today, well known technologies exist to scale applications to local networks but scaling to large networks with high latency is still a challenge. Load balancing at the middleware level allows more flexibility (in terms of granularity and distribution) than existing solutions based at lower system levels. However, it requires an execution infrastructure and mechanisms to be integr ...

15 Cluster reserves: a mechanism for resource management in cluster-based network servers 100%



Mohit Aron , Peter Druschel , Willy Zwaenepoel

**ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2000
 ACM SIGMETRICS international conference on Measurement and modeling of
 computer systems** June 2000

Volume 28 Issue 1

In network (e.g., Web) servers, it is often desirable to isolate the performance of different classes of requests from each other. That is, one seeks to achieve that a certain minimal proportion of server resources are available for a class of requests, independent of the load imposed by other requests. Recent work demonstrates how to achieve this performance isolation in servers consisting of a single, centralized node; however, achieving performance isolation in a distributed, cluster bas ...

16 Infinite parallel job allocation (extended abstract) 100%



Petra Berenbrink , Artur Czumaj , Tom Friedetzky , Nikita D. Vvedenskaya

**Proceedings of the twelfth annual ACM symposium on Parallel algorithms and
 architectures** July 2000

In recent years, the task of allocating jobs to servers has been studied with the "balls and bins" abstraction. Results in this area exploit the large decrease in maximum load that can be achieved by allowing each job (ball) a little freedom in choosing its destination server (bin). In this paper we examine an infinite and parallel allocation process (see [ABS98]) which is related to the "balls and bins" abstraction. The simple process can be used to model ...

17 Cluster resource management: Integrated resource management for cluster-based Internet services 100%



Kai Shen , Hong Tang , Tao Yang , Lingkun Chu

ACM SIGOPS Operating Systems Review December 2002

Volume 36 Issue SI

Client request rates for Internet services tend to be bursty and thus it is important to maintain efficient resource utilization under a wide range of load conditions. Network service clients typically seek services interactively and maintaining reasonable response time is often imperative for such services. In addition, providing differentiated service qualities and resource allocation to multiple service classes can also be desirable at times. This paper presents an integrated resource managem ...

18 Research sessions: potpourri: Workflow management with service quality guarantees 100%



Michael Gillmann , Gerhard Weikum , Wolfgang Wonner

Proceedings of the 2002 ACM SIGMOD international conference on Management of data June 2002

Workflow management systems (WFMS) that are geared for the orchestration of business processes across multiple organizations are complex distributed systems: they consist of multiple workflow engines, application servers, and communication middleware servers such as ORBs, where each of these server types can be replicated on multiple computers for scalability and availability. Finding an appropriate system configuration with guaranteed application-specific quality of service in terms of throughpu ...

19 Video Streaming 1: A Demand Adaptive and Locality Aware (DALA) 100%



streaming media server cluster architecture

Zihui Ge , Ping Ji , Prashant Shenoy

Proceedings of the 12th international workshop on Network and operating systems support for digital audio and video May 2002

The wide availability of broadband networking technologies such as cable modems and DSL coupled with the growing popularity of the Internet has led to a dramatic increase in the availability and the use of online streaming media. With the "last mile" network bandwidth no longer a constraint, the bottleneck for video streaming has been pushed closer to the server. Streaming high quality audio and video to a myriad of clients imposes significant resource demands on the server. In this work, we pro ...

20 Load distribution among replicated Web servers: a QoS-based 100%



approach

Marco Conti , Enrico Gregori , Fabio Panzieri

ACM SIGMETRICS Performance Evaluation Review March 2000

Volume 27 Issue 4

A dominant factor for the success of an Internet based Web service is the Quality of Service (QoS) perceived by its users. The principal QoS attributes these users perceive include those related to the service "responsiveness", i.e. the service availability and timeliness. In this paper, we argue that QoS can be provided by distributing the processing load among replicated Web servers, and that these servers can be geographically distributed across the Internet. In this context, we discuss strat ...

Results 1 - 20 of 200

short listing


Prev
Page

1 2 3 4 5 6 7 8 9 10


Next
Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[packet and relay <AND>((load and balance* and server and service))]**

Found **89** of **127,944** searched.

Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score** Blinder

Results 1 - 20 of 89 short listing



1 2 3 4 5



- 1 Dynamic task-based anycasting in mobile ad hoc networks 87%
 Prithwish Basu , Wang Ke , Thomas D. C. Little
Mobile Networks and Applications October 2003
 Volume 8 Issue 5

Mobile ad hoc networks (MANETs) have received significant attention in the recent past owing to the proliferation in the numbers of tetherless portable devices, and rapid growth in popularity of wireless networking. Most of the MANET research community has remained focused on developing lower layer mechanisms such as channel access and routing for making MANETs operational. However, little focus has been applied on higher layer issues, such as application modeling in dynamic MANET environments. ...

- 2 Fast detection of communication patterns in distributed executions 87%
 Thomas Kunz , Michiel F. H. Seuren
Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research November 1997

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

- 3 Stateful distributed interposition 85%
 John Reumann , Kang G. Shin
ACM Transactions on Computer Systems (TOCS) February 2004
 Volume 22 Issue 1

Interposition-based system enhancements for multitiered servers are difficult to build

because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex service ...

4 The effects of badly behaved routers on Internet congestion 84%



Kevin Curran , Derek Woods , Nadene McDermot , Colleen Bradley

International Journal of Network Management January 2003

Volume 13 Issue 1

From an exhaustive series of trace packets to a diverse set of destinations, our research has discovered that specific routers are the cause of bottlenecks in the Internet. We found that packets took the same route each time towards their destination. Our research has also found that over periods as large as seven days these routers continue to cause bottlenecks with no re-routing of packets to alleviate congestion. This research begs the question as to why these bottlenecks occur at the same place ...

5 Credit-based flow control for ATM networks: credit update protocol, adaptive credit allocation and statistical multiplexing 84%



H. T. Kung , Trevor Blackwell , Alan Chapman

ACM SIGCOMM Computer Communication Review , Proceedings of the conference on Communications architectures, protocols and applications October 1994

Volume 24 Issue 4

This paper presents three new results concerning credit-based flow control for ATM networks: (1) a simple and robust credit update protocol (CUP) suited for relatively inexpensive hardware/software implementation; (2) automatic adaptation of credit buffer allocation for virtual circuits (VCs) sharing the same buffer pool; (3) use of credit-based flow control to improve the effectiveness of statistical multiplexing in minimizing switch memory. These results have been substantiated by analysis ...

6 Network support for mobile multimedia using a self-adaptive distributed proxy 83%



Zhuoqing Morley Mao , Hoi-sheung Wilson So , Byunghoon Kang

Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video January 2001

Recent advancements in video and audio codec technologies~(e.g., RealVideo [18]) make multimedia streaming possible across a wide range of network conditions. With an increasing trend of ubiquitous connectivity, more and more areas have overlapping coverage of multiple wired and wireless networks. Because the best network service changes as the user moves, to provide good multimedia application performance, the service needs to adapt to user movement as well as network and computational resources ...

7 Protecting web servers from distributed denial of service attacks 83%



Frank Kargl , Joern Maier , Michael Weber

Proceedings of the tenth international conference on World Wide Web April 2001

8 Special feature: Report on a working session on security in wireless ad hoc networks 82%



Levente Buttyán , Jean-Pierre Hubaux

ACM SIGMOBILE Mobile Computing and Communications Review January 2003

Volume 7 Issue 1

- 9** Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer 82%



Michael J. Freedman , Robert Morris

Proceedings of the 9th ACM conference on Computer and communications security November 2002

Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a restricted topology in a way that adversaries cannot easily influence ...

- 10** Using name-based mappings to increase hit rates 82%



David G. Thaler , Chinya V. Ravishankar

IEEE/ACM Transactions on Networking (TON) February 1998
Volume 6 Issue 1

- 11** Cellular and hybrid networks: UCAN: a unified cellular and ad-hoc network architecture 82%



Haiyun Luo , Ramachandran Ramjee , Prasun Sinha , Li (Erran) Li , Songwu Lu

Proceedings of the 9th annual international conference on Mobile computing and networking September 2003

In third-generation (3G) wireless data networks, mobile users experiencing poor channel quality usually have low data-rate connections with the base-station. Providing service to low data-rate users is required for maintaining fairness, but at the cost of reducing the cell's aggregate throughput. In this paper, we propose the Unified Cellular and Ad-Hoc Network (UCAN) architecture for enhancing cell throughput, while maintaining fairness. In UCAN, a mobile client has both 3G cellular link and IE ...

- 12** Design and modelling of internode: a mobile provider provisioned VPN 82%



Francisco Barceló , Josep Paradells , Fofy Setaki , Monique Gibeaux

Mobile Networks and Applications February 2003
Volume 8 Issue 1

This paper presents the design and architecture of a mobile Provider Provisioned VPN (PPVPN) together with a performance evaluation oriented model that allows first estimates of the VPN set-up delay to be computed. At the same time, some consequences of the discussion can be applied to the design of the VPN configuration parameters. Many different technologies and protocols are used: access is supplied through GPRS or WaveLANs, IP mobility is supported by Mobile IP, and the VPN is based on the I ...

- 13** Virtual environments and HLA: A scalable architecture for supporting interactive games on the internet 82%



Wentong Cai , Percival Xavier , Stephen J. Turner , Bu-Sung Lee

Proceedings of the sixteenth workshop on Parallel and distributed simulation May 2002

This paper presents a scalable architecture for supporting large-scale interactive Internet games. In order to support a large number of participants and to divide the workload, the virtual world is divided into partitions. Each partition is then assigned to

a server. A client (i.e., a player or a participant) will join a server according to the position of the avatar it controls. Compared to a centralized architecture, this distributed client-server architecture is more scalable. In addition, c ...

14 Limits of wide-area thin-client computing

82%



Albert Lai , Jason Nieh

ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems June 2002

Volume 30 Issue 1

While many application service providers have proposed using thin-client computing to deliver computational services over the Internet, little work has been done to evaluate the effectiveness of thin-client computing in a wide-area network. To assess the potential of thin-client computing in the context of future commodity high-bandwidth Internet access, we have used a novel, non-invasive slow-motion benchmarking technique to evaluate the performance of several popular thin-client computing plat ...

15 Rethinking the design of the Internet: the end-to-end arguments vs. the

82%



brave new world

Marjory S. Blumenthal , David D. Clark

ACM Transactions on Internet Technology (TOIT) August 2001

Volume 1 Issue 1

This article looks at the Internet and the changing set of requirements for the Internet as it becomes more commercial, more oriented toward the consumer, and used for a wider set of purposes. We discuss a set of principles that have guided the design of the Internet, called the end-to-end arguments, and we conclude that there is a risk that the range of new requirements now emerging could have the consequence of compromising the Internet's original design principles. Were ...

16 ATM: retrospective on systems legacy: A retrospective view of ATM

80%



Charles Kalmanek

ACM SIGCOMM Computer Communication Review November 2002

Volume 32 Issue 5

ATM was the focus of active research and significant investment in the early to mid 1990's. This paper discusses several visions for ATM prevalent at the time, and analyzes how ATM evolved during this period. The paper also considers the implications of this history for current connection-oriented technologies, such as optical transport networks and MPLS.

17 A peering architecture for ubiquitous IP multicast streaming

80%



C. K. Yeo , B. S. Lee , M. H. Er

ACM SIGOPS Operating Systems Review July 2002

Volume 36 Issue 3

IP Multicast has been widely used for the distribution of real-time video. However, a large portion of the Internet is not multicast-enabled. Such networks rely on static relays and static tunnels to support video streaming. Being static, the streaming cannot readily adjust to the changing network conditions. This paper proposes an application level overlay using peering technology for seamless ubiquitous multimedia streaming across both multicast and unicast networks. The overlay comprises an a ...

18 Session A: Routing: On the impact of alternate path routing for load

80%



balancing in mobile ad hoc networks

Marc R. Pearlman , Zygmunt J. Haas , Peter Sholander , Siamak S. Tabrizi
Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing November 2000

Alternate path routing (APR) can provide load balancing and route failure protection by distributing traffic among a set of diverse paths. These benefits make APR appear to be an ideal candidate for the bandwidth limited and mobile ad-hoc networks. However, we find that APR's potential is not fully realized in ad-hoc networks because of route coupling resulting from the geographic proximity of candidate paths between common endpoints. In multiple channel networks, coupling occurs when paths shar ...

19 Papers: A survey of web caching schemes for the Internet

80%



Jia Wang

ACM SIGCOMM Computer Communication Review October 1999
Volume 29 Issue 5

The World Wide Web can be considered as a large distributed information system that provides access to shared data objects. As one of the most popular applications currently running on the Internet, the World Wide Web is of an exponential growth in size, which results in network congestion and server overloading. Web caching has been recognized as one of the effective schemes to alleviate the service bottleneck and reduce the network traffic, thereby minimize the user access latency. In this pap ...

20 Routing: On using the ad-hoc network model in cellular packet data networks

80%



Hung-Yun Hsieh , Raghupathy Sivakumar

Proceedings of the 3rd ACM international symposium on Mobile ad hoc networking & computing June 2002

While several approaches have been proposed in literature for improving the performance of wireless packet data networks, a recent class of approaches has focused on improving the underlying wireless network model itself. Several of such approaches have shown that using peer-to-peer communication, a mode of communication used typically in ad-hoc wireless networks, can result in performance improvement in terms of both throughput and energy consumption. However, the true impact of using the ad-ho ...

Results 1 - 20 of 89

short listing



Prev
Page

1 2 3 4 5



Next
Page

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.



[> home](#) [> about](#) [> feedback](#) [> login](#)

US Patent & Trademark Office



Try the *new* Portal design

Give us your opinion after using it.

Search Results

Search Results for: **[sessions<AND>((packet and relay <AND>((load and balance* and server and service))))]**

Found **43** of **127,944** searched.

Search within Results



[> Advanced Search](#)

[> Search Help/Tips](#)

Sort by: **Title** **Publication** **Publication Date** **Score** Binder

Results 1 - 20 of 43 **short listing**



1

2

3



- 1 Fast detection of communication patterns in distributed executions 87%
Thomas Kunz , Michiel F. H. Seuren
Proceedings of the 1997 conference of the Centre for Advanced Studies on Collaborative research November 1997

Understanding distributed applications is a tedious and difficult task. Visualizations based on process-time diagrams are often used to obtain a better understanding of the execution of the application. The visualization tool we use is Poet, an event tracer developed at the University of Waterloo. However, these diagrams are often very complex and do not provide the user with the desired overview of the application. In our experience, such tools display repeated occurrences of non-trivial commun ...

- 2 Network support for mobile multimedia using a self-adaptive distributed proxy 83%
Zhuoqing Morley Mao , Hoi-sheung Wilson So , Byunghoon Kang
Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video January 2001

Recent advancements in video and audio codec technologies~(e.g., RealV ideo [18] make multimedia streaming possible across a wide range of network conditions. With an increasing trend of ubiquitous connectivity, more and more areas have overlapping coverage of multiple wired and wireless networks. Because the best network service changes as the user moves, to provide good multimedia application performance, the service needs to adapt to user movement as well as network and computational res ...

- 3 Special feature: Report on a working session on security in wireless ad hoc networks 82%
Levente Buttyán , Jean-Pierre Hubaux
ACM SIGMOBILE Mobile Computing and Communications Review January 2003

Volume 7 Issue 1

- 4 Peer to peer networks: Tarzan: a peer-to-peer anonymizing network layer 82%



Michael J. Freedman , Robert Morris

Proceedings of the 9th ACM conference on Computer and communications security November 2002

Tarzan is a peer-to-peer anonymous IP network overlay. Because it provides IP service, Tarzan is general-purpose and transparent to applications. Organized as a decentralized peer-to-peer overlay, Tarzan is fault-tolerant, highly scalable, and easy to manage. Tarzan achieves its anonymity with layered encryption and multi-hop routing, much like a Chaumian mix. A message initiator chooses a path of peers pseudo-randomly through a restricted topology in a way that adversaries cannot easily influence ...

- 5 Using name-based mappings to increase hit rates 82%



David G. Thaler , Chinya V. Ravishankar

IEEE/ACM Transactions on Networking (TON) February 1998
Volume 6 Issue 1

- 6 Design and modelling of internode: a mobile provider provisioned VPN 82%



Francisco Barceló , Josep Paradells , Fofy Setaki , Monique Gibeaux

Mobile Networks and Applications February 2003
Volume 8 Issue 1

This paper presents the design and architecture of a mobile Provider Provisioned VPN (PPVPN) together with a performance evaluation oriented model that allows first estimates of the VPN set-up delay to be computed. At the same time, some consequences of the discussion can be applied to the design of the VPN configuration parameters. Many different technologies and protocols are used: access is supplied through GPRS or WaveLANs, IP mobility is supported by Mobile IP, and the VPN is based on the I ...

- 7 A peering architecture for ubiquitous IP multicast streaming 80%



C. K. Yeo , B. S. Lee , M. H. Er

ACM SIGOPS Operating Systems Review July 2002
Volume 36 Issue 3

IP Multicast has been widely used for the distribution of real-time video. However, a large portion of the Internet is not multicast-enabled. Such networks rely on static relays and static tunnels to support video streaming. Being static, the streaming cannot readily adjust to the changing network conditions. This paper proposes an application level overlay using peering technology for seamless ubiquitous multimedia streaming across both multicast and unicast networks. The overlay comprises an a ...

- 8 Session A: Routing: On the impact of alternate path routing for load balancing in mobile ad hoc networks 80%



Marc R. Pearlman , Zygmunt J. Haas , Peter Sholander , Siamak S. Tabrizi

Proceedings of the 1st ACM international symposium on Mobile ad hoc networking & computing November 2000

Alternate path routing (APR) can provide load balancing and route failure protection by distributing traffic among a set of diverse paths. These benefits make APR appear to be an ideal candidate for the bandwidth limited and mobile ad-hoc networks. However, we find that APR's potential is not fully realized in ad-hoc networks because of route

coupling resulting from the geographic proximity of candidate paths between common endpoints. In multiple channel networks, coupling occurs when paths shar ...

9 Session summaries from the 17th symposium on operating systems 80%



principle (SOSP'99)

Jay Lepreau , Eric Eide

ACM SIGOPS Operating Systems Review April 2000

Volume 34 Issue 2

10 Increasing the realibility of email services 80%



Joe Armstrong

Proceedings of the 2000 ACM symposium on Applied computing March 2000

11 A mechanism for supporting client migration in a shared window system 80%



Goopeel Chung , Prasun Dewan

Proceedings of the 9th annual ACM symposium on User interface software and technology November 1996

12 Stateful distributed interposition 80%



John Reumann , Kang G. Shin

ACM Transactions on Computer Systems (TOCS) February 2004

Volume 22 Issue 1

Interposition-based system enhancements for multitiered servers are difficult to build because important system context is typically lost at application and machine boundaries. For example, resource quotas and user identities do not propagate easily between cooperating services that execute on different hosts or that communicate with each other via intermediary services. Application-transparent system enhancement is difficult to achieve when such context information is obscured by complex servic ...

13 Application level performance: DNS performance and the effectiveness of caching 80%



Jaeyeon Jung , Emil Sit , Hari Balakrishnan , Robert Morris

Proceedings of the First ACM SIGCOMM Workshop on Internet Measurement Workshop November 2001

This paper presents a detailed analysis of traces of DNS and associated TCP traffic collected on the Internet links of the MIT Laboratory for Computer Science and the Korea Advanced Institute of Science and Technology (KAIST). The first part of the analysis details how clients at these institutions interact with the wide-area DNS system, focusing on performance and prevalence of failures. The second part evaluates the effectiveness of DNS caching. In the most recent MIT trace, 23% of lookups rece ...

14 A measurement analysis of Internet traffic over frame relay 80%



Judith L. Jerkins , John Monroe , Jonathan L. Wang

ACM SIGMETRICS Performance Evaluation Review September 1999

Volume 27 Issue 2

Various approaches have been proposed and implemented to relieve the congestion in the Public Switched Telephone Networks (PSTNs) induced by recent meteoric growth of Internet services. The Internet/Intranet Transport Service (IITS) offered by the Southwestern Bell Telephone (SWBT) Company provides an example of one such


implementation which off-loads the long-holding time data traffic from PSTNs to a packet technology (in this case Frame Relay). This paper describes analysis of 1997 IITS traffic ...

15 Manufacturing resource planning on a PC local area network 80%


 H. Clark Kee , Roy L. Post
ACM SIGAPL APL Quote Quad , Proceedings of the international conference on APL May 1986
 Volume 16 Issue 4

This paper details a large APL programming project of 12 man years. An integrated software system structured on the principles of MRP (manufacturing resource planning) was implemented by a Bristol-Myers in house team for use in a new manufacturing facility. The system applies off-the-shelf technology in innovative ways, using STSC APL*PLUS/PC as the only programming language, to build a very sophisticated application on IBM/PCs fully sharing data in a secure environment via the N ...

16 Protecting web servers from distributed denial of service attacks 77%


 Frank Kargl , Joern Maier , Michael Weber
Proceedings of the tenth international conference on World Wide Web April 2001

17 A multi-user framework supporting video-based avatars 77%

 Peter Quax , Tom Jehaes , Pieter Jorissen , Wim Lamotte
Proceedings of the 2nd workshop on Network and system support for games May 2003


In this paper we present our ongoing work in setting up a multi-user framework that supports video texturing on avatars, creating added value for both gaming and collaborative work applications. Limiting the data propagation and bandwidth usage is a key goal when targeting a scalable application that will be deployed on a general-purpose network such as the Internet. We therefore present a number of techniques that can be used in these circumstances, including increased client responsibilities a ...

18 Session 1: QoS's downfall: at the bottom, or not at all! 77%

 Jon Crowcroft , Steven Hand , Richard Mortier , Timothy Roscoe , Andrew Warfield
Proceedings of the ACM SIGCOMM workshop on Revisiting IP QoS: What have we learned, why do we care? August 2003

Quality of Service (QoS) has been touted as a technological requirement for many different networks at many different times. However, very few (if any) schemes for providing it have ever been successful, despite a huge amount of research in the area of QoS provision. In this position paper we analyze some of the reasons why so many QoS mechanisms have failed to be widely deployed. We suggest two factors in this failure: the timeliness of QoS mechanisms (they rarely arrive when they are needed), ...

19 DOS protection: Using graphic turing tests to counter automated DDoS attacks against web servers 77%

 William G. Morein , Angelos Stavrou , Debra L. Cook , Angelos D. Keromytis , Vishal Misra , Dan Rubenstein
Proceedings of the 10th ACM conference on Computer and communication security October 2003

We present WebSOS, a novel overlay-based architecture that provides guaranteed access to a web server that is targeted by a denial of service (DoS) attack. Our

approach exploits two key characteristics of the web environment: its design around a human-centric interface, and the extensibility inherent in many browsers through downloadable "applets." We guarantee access to a web server for a large number of *previously unknown* users, without requiring pre-existing trust relationships between ...

20 Overlay & peer-to-peer networks: SplitStream: high-bandwidth

77%



multicast in cooperative environments

Miguel Castro , Peter Druschel , Anne-Marie Kermarrec , Animesh Nandi , Antony Rowstron , Atul Singh

Proceedings of the nineteenth ACM symposium on Operating systems principles

October 2003

In tree-based multicast systems, a relatively small number of interior nodes carry the load of forwarding multicast messages. This works well when the interior nodes are highly-available, dedicated infrastructure routers but it poses a problem for application-level multicast in peer-to-peer systems. SplitStream addresses this problem by striping the content across a forest of interior-node-disjoint multicast trees that distributes the forwarding load among all participating peers. For example, i ...

Results 1 - 20 of 43 **short listing**



1

2

3



The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE



Membership Publications/Services Standards Conferences Careers/Jobs

IEEE Xplore®
 RELEASE 1.6

 Welcome
 United States Patent and Trademark Office

[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)
[Quick Links](#)

Welcome to IEEE Xplore®

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

Your search matched **24** of **1011253** documents.A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.**Refine This Search:**

You may refine your search by editing the current search expression or enter a new one in the text box.

load and balance* and server and service

Search☒ Check to search within this result set**Results Key:****JNL** = Journal or Magazine **CNF** = Conference **STD** = Standard**1 Managing services in distributed systems by integrating trading and balancing***Thissen, D.; Neukirchen, H.;*Computers and Communications, 2000. Proceedings. ISCC 2000. Fifth IEEE Symposium on , 3-6 July 2000
Pages:641 - 646
[\[Abstract\]](#) [\[PDF Full-Text \(556 KB\)\]](#) **IEEE CNF**
2 Adaptive optimal load balancing in a nonhomogeneous multiserver system with a central job scheduler*Bonomi, F.; Kumar, A.;*Computers, IEEE Transactions on , Volume: 39 , Issue: 10 , Oct. 1990
Pages:1232 - 1250
[\[Abstract\]](#) [\[PDF Full-Text \(1212 KB\)\]](#) **IEEE JNL**
3 An intelligent load distribution system for CORBA-compliant distributed environments*Damiani, E.;*Fuzzy Systems Conference Proceedings, 1999. FUZZ-IEEE '99. 1999 IEEE International , Volume: 1 , 22-25 Aug. 1999
Pages:332 - 336 vol.1
[\[Abstract\]](#) [\[PDF Full-Text \(400 KB\)\]](#) **IEEE CNF**
4 Video-on-demand server system design with random early migration*Yinqing Zhao; Kuo, C.-C.J.;*

Circuits and Systems, 2003. ISCAS '03. Proceedings of the 2003 International Symposium on , Volume: 2 , 25-28 May 2003

Pages:II-640 - II-643 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(348 KB\)\]](#) IEEE CNF

5 Speculative data dissemination and service to reduce server load, network traffic and service time in distributed information systems

Bestavros, A.;

Data Engineering, 1996. Proceedings of the Twelfth International Conference on , 26 Feb.-1 March 1996

Pages:180 - 187

[\[Abstract\]](#) [\[PDF Full-Text \(844 KB\)\]](#) IEEE CNF

6 WebWave: globally load balanced fully distributed caching of hot published documents

Heddaya, A.; Mirdad, A.;

Distributed Computing Systems, 1997., Proceedings of the 17th International Conference on , 27-30 May 1997

Pages:160 - 168

[\[Abstract\]](#) [\[PDF Full-Text \(840 KB\)\]](#) IEEE CNF

7 Load testing of Web sites

Menasce, D.A.;

Internet Computing, IEEE , Volume: 6 , Issue: 4 , July-Aug. 2002

Pages:70 - 74

[\[Abstract\]](#) [\[PDF Full-Text \(359 KB\)\]](#) IEEE JNL

8 Data distribution algorithms for load balanced fault-tolerant Web applications

Narendran, B.; Rangarajan, S.; Yajnik, S.;

Reliable Distributed Systems, 1997. Proceedings., The Sixteenth Symposium on , 22-24 Oct. 1997

Pages:97 - 106

[\[Abstract\]](#) [\[PDF Full-Text \(976 KB\)\]](#) IEEE CNF

9 Right-sizing computer resources for re-engineered processes

Weldon, D.E.;

Aerospace and Electronics Conference, 1996. NAECON 1996., Proceedings of IEEE 1996 National , Volume: 2 , 20-23 May 1996

Pages:825 - 828 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(340 KB\)\]](#) IEEE CNF

10 Optimality of weighted least squares load balancing

Bonomi, F.; Kumar, A.;

Decision and Control, 1988., Proceedings of the 27th IEEE Conference on , 7-Dec. 1988

Pages:1480 - 1485 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(440 KB\)\]](#) IEEE CNF

11 Integrated fault-tolerant multicast and anycast routing algorithms*Jia, W.; Xu, G.; Zhao, W.;*Computers and Digital Techniques, IEE Proceedings- , Volume: 147 , Issue: 4
2000

Pages:266 - 274

[\[Abstract\]](#) [\[PDF Full-Text \(948 KB\)\]](#) [IEEE JNL](#)**12 Parallel cache management protocol and algorithm for cooperative servers***Jaeyong Lim; Rajasekaran, S.;*Communications, 2003. ICC '03. IEEE International Conference on , Volume:
2 , 11-15 May 2003

Pages:918 - 922 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(312 KB\)\]](#) [IEEE CNF](#)**13 Dynamic scheduling of Web server cluster***Suntae Hwang; Naksoo Jung;*Parallel and Distributed Systems, 2002. Proceedings. Ninth International
Conference on , 17-20 Dec. 2002

Pages:563 - 568

[\[Abstract\]](#) [\[PDF Full-Text \(799 KB\)\]](#) [IEEE CNF](#)**14 On the placement of Web server replicas***Lili Qiu; Padmanabhan, V.N.; Voelker, G.M.;*INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and
Communications Societies. Proceedings. IEEE , Volume: 3 , 22-26 April 2001

Pages:1587 - 1596 vol.3

[\[Abstract\]](#) [\[PDF Full-Text \(524 KB\)\]](#) [IEEE CNF](#)**15 A fuzzy approach to load balancing in a distributed object computing network***Lap-Sun Cheung;*Cluster Computing and the Grid, 2001. Proceedings. First IEEE/ACM Internati
Symposium on , 15-18 May 2001

Pages:694 - 699

[\[Abstract\]](#) [\[PDF Full-Text \(492 KB\)\]](#) [IEEE CNF](#)[1](#) [2](#) [Next](#)

[Home](#) | [Log-out](#) | [Journals](#) | [Conference Proceedings](#) | [Standards](#) | [Search by Author](#) | [Basic Search](#) | [Advanced Search](#) | [Join IEEE](#) | [Web Account](#) |
[New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online](#)
[Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

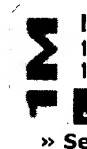
Copyright © 2004 IEEE — All rights reserved

IEEE HOME | SEARCH IEEE | SHOP | WEB ACCOUNT | CONTACT IEEE



Membership Publications/Services Standards Conferences Careers/Jobs

IEEE Xplore®
 RELEASE 1.6

 Welcome
 United States Patent and Trademark Office

[Help](#) [FAQ](#) [Terms](#) [IEEE Peer Review](#)
[Quick Links](#)

Welcome to IEEE Xplore®

- ☐ Home
- ☐ What Can I Access?
- ☐ Log-out

Tables of Contents

- ☐ Journals & Magazines
- ☐ Conference Proceedings
- ☐ Standards

Search

- ☐ By Author
- ☐ Basic
- ☐ Advanced

Member Services

- ☐ Join IEEE
- ☐ Establish IEEE Web Account
- ☐ Access the IEEE Member Digital Library

Your search matched **24** of **1011253** documents.A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.**Refine This Search:**

You may refine your search by editing the current search expression or enter a new one in the text box.

load and balance* and server and service

Search☐ Check to search within this result set**Results Key:****JNL** = Journal or Magazine **CNF** = Conference **STD** = Standard**16 Storage server for processing video transactions through sequential access in multimedia DBMS***Keun Hyung Kim; Seog Park;*Database Applications in Non-Traditional Environments, 1999. (DANTE '99) Proceedings. 1999 International Symposium on , 1999
Pages:168 - 175
[\[Abstract\]](#) [\[PDF Full-Text \(312 KB\)\]](#) **IEEE CNF**
17 Operating firewalls outside the LAN perimeter*Smith, R.N.; Bhattacharya, S.;*Performance, Computing and Communications Conference, 1999. IPCCC '99. International , 10-12 Feb. 1999
Pages:493 - 498
[\[Abstract\]](#) [\[PDF Full-Text \(508 KB\)\]](#) **IEEE CNF**
18 BALANCE-a flexible parallel load balancing system for heterogeneous computing systems and networks*Chi-Chung Hui; Chanson, S.T.; Pui-Man Chui; Ka-Ming Lau;*

INFOCOM '96. Fifteenth Annual Joint Conference of the IEEE Computer Society Networking the Next Generation. Proceedings IEEE , Volume: 2 , 24-28 March

Pages:896 - 903 vol.2

[\[Abstract\]](#) [\[PDF Full-Text \(876 KB\)\]](#) **IEEE CNF**
19 Analysis of token ring networks with both finite and infinite buffer capacity stations*Deng, R.H.; Chiew, W.C.L.; Huang, K.-T.;*

INFOCOM '91. Proceedings. Tenth Annual Joint Conference of the IEEE Computer and Communications Societies. Networking in the 90s. IEEE , 7-11 April 1991
Pages:202 - 210 vol.1

[\[Abstract\]](#) [\[PDF Full-Text \(572 KB\)\]](#) [IEEE CNF](#)

20 Replicating the contents of a WWW multimedia repository to minimize download time

Loukopoulos, T.; Ahmad, I.;

Parallel and Distributed Processing Symposium, 2000. IPDPS 2000. Proceedings. 14th International , 1-5 May 2000

Pages:500 - 505

[\[Abstract\]](#) [\[PDF Full-Text \(112 KB\)\]](#) [IEEE CNF](#)

21 Load balancing a cluster of web servers: using distributed packet rewriting

Aversa, L.; Bestavros, A.;

Performance, Computing, and Communications Conference, 2000. IPCCC '00. Conference Proceeding of the IEEE International , 20-22 Feb. 2000

Pages:24 - 29

[\[Abstract\]](#) [\[PDF Full-Text \(864 KB\)\]](#) [IEEE CNF](#)

22 Load balancing and hot spot relief for hash routing among a collection of proxy caches

Kun-Lung Wu; Yu, P.S.;

Distributed Computing Systems, 1999. Proceedings. 19th IEEE International Conference on , 31 May-4 June 1999

Pages:536 - 543

[\[Abstract\]](#) [\[PDF Full-Text \(716 KB\)\]](#) [IEEE CNF](#)

23 Walrus-a low latency, high throughput Web service using Internet replication

Amir, Y.; Shaw, D.;

Electronic Commerce and Web-based Applications/Middleware, 1999. Proceedings. 19th IEEE International Conference on Distributed Computing Systems Workshop on , 31 May-4 June 1999

Pages:31 - 40

[\[Abstract\]](#) [\[PDF Full-Text \(144 KB\)\]](#) [IEEE CNF](#)

24 Implementation of a parallel and distributed MAKE on NFS with GAFolliot, B.;

Computers and Communications, 1990. Conference Proceedings., Ninth Annual International Phoenix Conference on , 21-23 March 1990

Pages:871

[\[Abstract\]](#) [\[PDF Full-Text \(104 KB\)\]](#) [IEEE CNF](#)

[Prev](#) [1](#) [2](#)

[New this week](#) | [OPAC Linking Information](#) | [Your Feedback](#) | [Technical Support](#) | [Email Alerting](#) | [No Robots Please](#) | [Release Notes](#) | [IEEE Online Publications](#) | [Help](#) | [FAQ](#) | [Terms](#) | [Back to Top](#)

Copyright © 2004 IEEE — All rights reserved